

Appl. No. 10/735,428
Filing Date 12/12/2003

RECEIVED
CENTRAL FAX CENTER
MAY 10 2006

JP920030192US1

IN THE CLAIMS

Please amend the claims as follows:

1. (currently amended) A method for selecting an input/output scheduler in a computing system having a plurality of input/output schedulers, said method comprising the steps of:

a) mapping, by an operating system kernel, each of said plurality of operating system input/output schedulers against a corresponding desired respective sets of heuristics, wherein a set of heuristic variables for characterizing performance states of the computing system include:

a first variable for a number of I/O job requests as a proportion of a total number of processes requesting I/O jobs,

a second variable for a number of read operations as a proportion of total number of read and write operations, and

a third variable for an average disk seek distance of submitted job requests,

and wherein the heuristic sets include:

a first one of the heuristic sets defining a first performance state of the computing system, wherein a certain predefined high proportion of read operations is exceeded for the second variable and a certain predefined average seek distance is not exceeded for the third variable,

a second one of the heuristic sets defining a second performance state of the computing system, wherein a certain predefined high level of I/O job requests per process is exceeded for the first variable, a certain predefined low proportion of read operations is not exceeded for the second variable, and the certain predefined average seek distance is not exceeded for the third variable, and

a third one of the heuristic sets defining a third performance state of the computing system, wherein the certain predefined high level of I/O job requests per process is exceeded for the first variable and the certain predefined low

Appl. No. 10/735,428
Filing Date 12/12/2003

JP920030192US1

proportion of read operations is exceeded but the certain predefined high proportion of read operations is not exceeded for the second variable,
wherein the operating system input/output schedulers include first, second and third input/output schedulers, and the mapping includes the operating system kernel:
associating the first one of the heuristic sets with the first scheduler, the second one of the heuristic sets with the second scheduler and the third one of the heuristic sets with the third scheduler;
b) monitoring heuristics relating to job requests performed in performance of said computing system, including a thread of the operating system kernel monitoring performance values of the heuristic variables; and
c) selecting one of the input/output schedulers by the operating system kernel thread, wherein the selecting is responsive to the operating system kernel thread determining that a performance state of the computing system is such that values of whether said monitored heuristics variables match one any of said desired first, second, or third sets of heuristics to select one of said plurality of input/output schedulers.

2. (currently amended) The method according to claim 1, wherein the operating system input/output schedulers include a fourth input/output scheduler and the method includes comprising the further step of:

selecting the fourth a-default input/output scheduler, from said plurality of input/output schedulers, when wherein the selecting is responsive to the operating system kernel thread determining that a performance state of the computing system is such that values of said monitored heuristics variables do not match any of said first, second or third desired sets of heuristics.

3. (currently amended) The method according to claim 1, wherein the set of heuristic variables include said monitored heuristics are selected from the group of heuristics consisting of number of read requests, number of write requests, proportion of read requests to write requests, average time for a process to submit successive job requests, input/output throughput, and disk utilization.

Appl. No. 10/735,428
Filing Date 12/12/2003

JP920030192US1

4. (currently amended) The method according to claim 1, wherein including:
scheduling input/output operations responsive to an anticipatory scheduling
algorithm by the first input/output scheduler;
scheduling input/output operations responsive to a first in first out algorithm by
the second input/output scheduler; and
scheduling input/output operations responsive to a fairness queue algorithm by the
third input/output scheduler. ~~at least one of said plurality of input/output schedulers is~~
~~selected from the group of schedulers consisting of First In First Out (FIFO), Shortest~~
~~Positioning Time First (SPTF), Anticipatory Scheduler (AS), Deadline scheduler, and~~
~~Fairness Queue scheduler.~~

5 – 7. (canceled)

8. (currently amended) A computing system for selecting an input/output scheduler, said computing system comprising:
at least one application;
a plurality of first, second, and third input/output schedulers;
a mapping of said plurality of input/output schedulers against a corresponding
desired first, second, and third sets of operating heuristics; wherein the heuristic sets
include first, second, and third heuristic variables, the first variable being for a number of
I/O job requests as a proportion of a total number of processes requesting I/O jobs, the
second variable being for a number of read operations as a proportion of total number of
read and write operations, and the third variable being for an average disk seek distance
of submitted job requests, and wherein a first one of the heuristic sets defines a first
performance state of the computing system, wherein a certain predefined high proportion
of read operations is exceeded for the second variable and a certain predefined average
seek distance is not exceeded for the third variable, a second one of the heuristic sets
defines a second performance state of the computing system, wherein a certain
predefined high level of I/O job requests per process is exceeded for the first variable, a
certain predefined low proportion of read operations is not exceeded for the second
variable, and the certain predefined average seek distance is not exceeded for the third

Appl. No. 10/735,428
Filing Date 12/12/2003

JP920030192US1

variable, and a third one of the heuristic sets defines a third performance state of the computing system, wherein the certain predefined high level of I/O job requests per process is exceeded for the first variable and the certain predefined low proportion of read operations is exceeded but the certain predefined high proportion of read operations is not exceeded for the second variable, and wherein the mapping includes associations among respective first, second and third ones of the heuristic sets and first, second and third ones of the input/output schedulers; and

an operating system kernel for gathering and analysing heuristics relating to job requests submitted to said operating system kernel by said at least one application, said operating system kernel selecting one of said first, second, and third plurality of input/output schedulers based on said analysed heuristics and said desired sets of operating heuristics in said table, responsive to the operating system kernel thread determining that a performance state of the computing system is such that values of said heuristic variables match one of said first, second or third sets of heuristics.

9. (currently amended) The computing system according to claim 8, wherein the set of heuristic variables includes said analysed heuristics are selected from the group of heuristics consisting of number of read requests, number of write requests, proportion of read requests to write requests, average time for a process to submit successive job requests, input/output throughput, and disk utilization.

10. (currently amended) The computing system according to claim 8, wherein the first input/output scheduler is operable to schedule input/output operations responsive to an anticipatory scheduling algorithm, the second input/output scheduler is operable to schedule input/output operations responsive to a first in first out algorithm and the third input/output scheduler is operable to schedule input/output operations responsive to a fairness queue algorithm, at least one of said plurality of input/output schedulers is selected from the group of schedulers consisting of First In First Out (FIFO), Shortest Positioning Time First (SPTF), Anticipatory Scheduler (AS), Deadline scheduler, and Fairness Queue scheduler.

Appl. No. 10/735,428
Filing Date 12/12/2003

JP920030192US1

11. (original) An operating system kernel in a computing system, said operating system kernel comprising:

means for maintaining a predetermined set of values associated with a plurality of input/output schedulers;

means for monitoring heuristics associated with jobs performed in said computing system; and

means for comparing said monitored heuristics with said predetermined set of values to select one of said plurality of input/output schedulers.

12. (original) The operating system kernel according to claim 11, further comprising:
means for activating said selected input/output scheduler.

13. (original) The operating system kernel according to claim 11, wherein one of said plurality of input/output schedulers is a default scheduler.

14. (currently amended) The operating system kernel according to claim 11, wherein said predetermined set of values are selected from a the group of heuristics consisting of number of read requests, number of write requests, proportion of read requests to write requests, average time for a process to submit successive job requests, input/output throughput, and disk utilization.;

15. (original) The operating system kernel according to claim 11, wherein said activating means is a daemon.

16. (currently amended) A computing system having a plurality of available input/output schedulers, said computing system comprising:

an operating system module;

an input/output scheduling module having an active input/output scheduler selected from said plurality of input/output schedulers;

a hardware device drivers module for executing job requests received from said operating system module via said input/output scheduling module;

Appl. No. 10/735,428
Filing Date 12/12/2003

JP920030192US1

a heuristics module in an operating system kernel of the computing system, the heuristics module being operable for analysing information returned from said hardware device drivers module relating to said executed job requests; and

a switch module in an operating system kernel of the computing system, the switch module being operable for comparing said analysed information with a predetermined, mutually exclusive sets of heuristic values to identify one of the mutually exclusive sets of heuristic values select a preferred one of said plurality of input/output schedulers; and a kernel thread for switching said active input/output scheduler in said input/output scheduling module to a selected one of said preferred input/output schedulers, wherein the selected one of the input/output schedulers corresponds to the identified one of the mutually exclusive sets of heuristic values.

17. (currently amended) A computer program product comprising a computer readable medium having a computer program recorded therein for selecting an input/output scheduler in a computing system having a plurality of input/output schedulers, said computer program comprising:

computer program code means for mapping, by an operating system kernel, each of said plurality of input/output schedulers against a corresponding desired respective sets of heuristics, wherein a set of heuristic variables for characterizing performance states of the computing system includes:

a first variable for a number of I/O job requests as a proportion of a total number of processes requesting I/O jobs,

Appl. No. 10/735,428
Filing Date 12/12/2003

JP920030192US1

a second variable for a number of read operations as a proportion of total number of read and write operations, and

a third variable for an average disk seek distance of submitted job requests, and

wherein the computer program code means for mapping includes:

computer program code means for a first one of the heuristic sets defining a first performance state of the computer system, wherein a certain predefined high proportion of read operations is exceeded for the second variable and a certain predefined average seek distance is not exceeded for the third variable,

computer program code means for a second one of the heuristic sets defining a second performance state of the computer system, wherein a certain predefined high level of I/O job requests per process is exceeded for the first variable, a certain predefined low proportion of read operations is not exceeded for the second variable, and the certain predefined average seek distance is not exceeded for the third variable, and

computer program code means for a third one of the heuristic sets defining a third performance state of the computer system, wherein the certain predefined high level of I/O job requests per process is exceeded for the first variable and the certain predefined low proportion of read operations is exceeded but the certain predefined high proportion of read operations is not exceeded for the second variable;

computer program code means for monitoring heuristics relating to job requests performed in said computing system, wherein the computer program code means for monitoring includes computer program code means for a thread of the operating system kernel monitoring performance values of the heuristic variables; and

computer program code means for the operating system kernel thread selecting one of the input/output schedulers, wherein the selecting is responsive to the operating system kernel thread determining that a performance state of the computing system is such that values of whether said monitored heuristics variables

Appl. No. 10/735,428
Filing Date 12/12/2003

JP920030192US1

match one any of said desired first, second, or third sets of heuristics, to select one of said plurality of input/output schedulers.

18. (canceled)
19. (new) The computer program product of claim 17, comprising:
computer program code means for the first input/output scheduler scheduling input/output operations responsive to an anticipatory scheduling algorithm;
computer program code means for the second input/output scheduler scheduling input/output operations responsive to a first in first out algorithm; and
computer program code means for the third input/output scheduler scheduling input/output operations responsive to a fairness queue algorithm.